DOCUMENT RESUME

ED 222 187 IR 010 421

AUTHOR Mann, William C.; And Others

TITLE Text Generation: The State of the Art and the Literature.

INSTITUTION University of Southern California, Marina del Rey. Information Sciences Inst.

SPONS AGENCY Air Force Office of Scientific Research, Washington, D.C.
REPORT NO ISI/RR-81-101
PUB DATE Dec 81
NOTE 25p.

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS *Artificial Intelligence; *Computational Linguistics; Computer Programs; *Design Requirements; Discourse Analysis; Generative Grammar; *Models; *Research Needs
IDENTIFIERS *Natural Language Processing; *Text Generation; Text Structure

ABSTRACT

This report comprises two documents which describe the state of the art of computer generation of natural language text. Both were prepared by a panel of individuals who are active in research on text generation. The first document assesses the techniques now available for use in systems design, covering all of the technical methods by which computer programs create and present English text in their outputs, and identifies four kinds of technical developments which will shape the art in the coming decade: linguistically justified grammars, knowledge representation methods, models of the reader, and models of discourse. A comprehensive bibliography on text generation, the second document lists published works on the generation of natural language text by computer programs and some prior noncomputational work that has been used as a basis for such computer programs. Seven brief descriptions are included for research in progress or research recently completed but not yet described comprehensively in any publication. (Author/LMM)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED222187

William C. Mann
USC Information Sciences Institute
Madeline Bates
Bolt Beranek and Newman
Barbara J. Grosz
SRI International
David D. McDonald
University of Massachusetts at Amherst
Kathleen R. McKeown
University of Pennsylvania
William R. Swartout
USC Information Sciences Institute

U.S. DEPARTMENT OF EDUCATION
NATIONAL INSTITUTE OF EDUCATION
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

- This document has been reproduced as received from the person or organization originating it.
  Minor changes have been made to improve reproduction quality.
- Points of view or opinions stated in this document do not necessarily represent official NIE position or policy.

ISI/RR-81-101 and
University of Pennsylvania
MS-CIS-81-9
December 1981

Text Generation:
The State of the Art and the Literature

LR010421

INFORMATION SCIENCES INSTITUTE
UNIVERSITY OF SOUTHERN CALIFORNIA
4676 Admiralty Way/Marina del Rey/California 90291
(213) 822-1511

The views and conclusions contained in this report are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research of the U.S. Government.

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ISI/RR-81-101	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Text Generation: The State of the Art and the Literature		5. TYPE OF REPORT & PERIOD COVERED Research Report 6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) William C. Mann, Madeline Bates, Barbara J. Grosz, David D. McDonald, Kathleen R. McKeown, William R. Swartout		8. CONTRACT OR GRANT NUMBER(s) (OVER)
9. PERFORMING ORGANIZATION NAME AND ADDRESS USC/Information Sciences Institute 4676 Admiralty Way Marina del Rey, CA 90291		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Mathematical and Information Sciences Air Force Office of Scientific Research Building 410 Bolling Air Force Base Washington, D.C. 20332		12. REPORT DATE December 1981 13. NUMBER OF PAGES 26
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

This document is approved for public release and sale; distribution is unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

.....

18. SUPPLEMENTARY NOTES

.....

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

bibliography, computer generation of text, linguistically justified grammars, knowledge representation methods, models of the reader, models of discourse, natural language, text generation

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report consists of two documents describing the state of the art of computer generation of natural language text. Both were prepared by a panel of individuals who are active in research on text generation. The first document assesses the state of the art, identifying four kinds of technical developments which will shape the art in the coming decade: linguistically justified grammars, knowledge representation methods, models of the reader, and models of discourse. The second document is a comprehensive bibliography on text generation, the first of its kind. In addition to citations of documents, it includes descriptions of ongoing research efforts.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

8. CONTRACT OR GRANT NUMBER(s)

MCS76-07332
MCS79-08401
MCS81-07290
F49620-79-C-0181
N00039-80-C-0645
N00039-80-C-575

20. ABSTRACT (continued)

The authors were convened as a panel by the Naval Research Laboratory in conjunction with the 1981 Conference of the Association for Computational Linguistics.

PREFACE

This report consists of two documents describing the state of the art of computer generation of natural language text. Both were prepared by a panel of individuals who are active in research on text generation. The first document assesses the state of the art, identifying four kinds of technical developments which will shape the art in the coming decade: linguistically justified grammars, knowledge representation methods, models of the reader, and models of discourse. The second document is a comprehensive bibliography on text generation, the first of its kind. In addition to citations of documents, it includes descriptions of ongoing research efforts.

The authors were convened as a panel by the Naval Research Laboratory in conjunction with the 1981 Conference of the Association for Computational Linguistics.

PART 1. THE STATE OF THE ART IN TEXT GENERATION: SUMMARY AND PROJECTION

Contents

1. ASSESSING TEXT GENERATION TECHNOLOGY 1

2. WHAT TECHNIQUES ARE NOW AVAILABLE FOR USE IN SYSTEM DESIGNS? 1

3. BASIC COMPONENTS FOR A TEXT GENERATION FACILITY 2

4. LINGUISTICALLY JUSTIFIED GRAMMARS FOR TEXT GENERATION 3

5. KNOWLEDGE REPRESENTATION FORMALISMS 4

6. MODELS OF THE READER 4

7. MODELS OF DISCOURSE 5

8. RELATING THE BASIC COMPONENTS TO TEXT GENERATION 5
- 8.1. The Text Quality Limitations of Grammars 5
- 8.2. The Knowledge Representation 6
- 8.3. Models of the Reader or System User 6
- 8.4. Models of Discourse 6

9. DESIGNING IN 1986 FOR PRACTICAL TEXT GENERATION 7

10. PRESENT RESEARCH STATUS 7

1. ASSESSING TEXT GENERATION TECHNOLOGY

Our goal here is to assess the state of the art of text generation for two purposes: to help people who intend to apply that art in the near future and to aid in the design or selection of appropriate research.

This assessment covers all of the technical methods by which computer programs create and present English text in their outputs. (For simplicity we always call the output language English.) Because text generation has not always been taken seriously from a technical point of view, it has been actively pursued only recently as a topic in artificial intelligence. As a result of this late start, much of the technology available for application today is still rather superficial. However, text generation is now such an active research topic that this superficial technology will soon be surpassed.[1]

2. WHAT TECHNIQUES ARE NOW AVAILABLE FOR USE IN SYSTEM DESIGNS?

Two kinds of practical text generation techniques are already in general use and fairly well understood. The first is displaying previously prepared text (or ***canned text***), and the second is producing text by ***direct translation of knowledge structures.***

The simplest and most commonly used way to have a computer system produce text is for the implementors of the system to figure out in advance what sorts of English output will be required and then store it as text strings. The computer merely displays the text that has been stored. (For example, almost all error messages are produced in this way.) It is relatively easy to have a program produce English in this way, and the text can be complex and elegantly written if desired. Unfortunately, because the text strings can be changed independently of any knowledge structures the program might use, there is no guarantee of consistency between what the program does and what it says it does. Another problem with canned text is that all questions and answers must be anticipated in advance; for large systems, that may prove to be impossible. Finally, since one text string looks like any other as far as the computer is concerned, the computer program cannot easily have a conceptual model of what it is saying. This means that one should not expect to see much closure: satisfying 100 needs for text will not make the second 100 much easier.

Another approach to providing English output produces text by translating knowledge structures of the program directly to English. This method overcomes many of the problems with canned text, while introducing some of its own. Since the structures being transformed (translated) are the same

---

[1] This part of the report does not attempt to index the literature on text generation. However, the second part of this report contains an extensive bibliography on the subject.

ones used in the program's reasoning process, consistency can be assured. Closure can be realized because transformations are written to handle large classes of knowledge structures. However, since the transformations performed are usually relatively simple, the quality of the text depends to a great degree on how the knowledge is structured. If the text is to be understandable, the knowledge used by the program must be structured so that it is readily understood. Finally, systems employing this technique typically have had very little linguistic knowledge, so they have produced text that is verbose, stilted, and redundant, although readable.

Practical, near-term applications of text generation will share certain characteristics:

1. They require short texts: one to three sentences.
2. They have well-elaborated program data structures corresponding in fairly simple ways to the desired texts.
3. The important knowledge can be represented well with present techniques; it does not involve the difficulties listed in section 5.
4. Limited fluency of output is acceptable.

Some so-called "expert systems" that explain their reasoning in English have these characteristics.

We believe that text-producing systems of the future will continue to include processes that produce text by translating knowledge structures. However, they will be integrated with other processes that use extensive linguistic knowledge, a discourse model, a model of the reader, and enhanced knowledge representations.

Because of the limited capabilities of present techniques, a new project aiming to produce a benchmark application program in the text generation area ***would currently be counterproductive,*** since it would produce little or no transferable technology and would detract from the community's ability to make progress on the general problem.

3. BASIC COMPONENTS FOR A TEXT GENERATION FACILITY

How can the very limited capabilities now available be developed into fluent, powerful text generation methods that are easily applied to new tasks? The next few sections describe the kinds of methods that are needed and are being developed.

The underlying model presumed here, which present research is moving toward, has the following characteristics:

1. Responsibility for text generation is in a text generation module rather than being scattered at the points of use.

2. A major portion of the text generation module is portable and is developed cumulatively through many systems. The portable components include a grammar that encodes general knowledge of English and processes that handle linguistic, task-independent information.

We feel strongly that a competent text generation facility must have the following four identifiable components, and that ***limitations on these will be limitations on the overall state of the art for the foreseeable future:***

1. A comprehensive, linguistically justified grammar.
2. A knowledge-representation formalism that can encode diverse kinds of information.
3. A model of the intended reader of the text.
4. A model of discourse structure and control.

Each of these draws on existing noncomputational precedents, and each requires some special adaptation to the text generation task.

Below we describe each of these basic components in a form that it might achieve in five to ten years of research. (These descriptions are followed by a projection of the practical alternatives available to system designers five years hence.)

4. LINGUISTICALLY JUSTIFIED GRAMMARS FOR TEXT GENERATION

Grammars are ordinarily developed by linguists over periods of ten to twenty or more years, in departments of linguistics. The best ones may be written by a single individual, but they reflect the ideas of dozens or hundreds of people who have contributed to refining particular forms.

Present practice in linguistics emphasizes carefully reasoned development of small fragments of grammars. Hence comprehensive, linguistically justified grammars, the sort we need, are very rare.

Several linguistic traditions (some associated with computation and some not) are particularly likely to produce suitably refined, comprehensive grammars for text generators. They are:

1. The systemic tradition, founded by Michael A. K. Halliday around 1961.
2. The transformational tradition, decisively articulated by Noam Chomsky starting in 1958.
3. The Generalized Phrase Structure Grammar tradition, currently associated with Gerald Gazdar.
4. The ATN tradition, begun by Bill Woods and now being developed by him and many others.
5. The LSP tradition, developed so far mainly under the direction of Naomi Sager.
6. The Diamond (or Diagram) grammar by Jane Robinson.

Grammars do not appear in computers without extensive effort. Most linguists are not interested in providing or seeing the level of detail and precise definition needed for effective computational use. There are enormous social barriers between the source of these grammars (linguists) and their potential users. It will be necessary to sponsor text generation research projects with linguists on the staff: projects staffed entirely by computer people can be expected to yield only short-term results.

5. KNOWLEDGE REPRESENTATION FORMALISMS

Text generation programs cannot improve much on the knowledge they are given. The notation for knowledge must already contain appropriate abstractions in an easily accessible form. Today's notations are relatively good at representing logical formulas and deductive necessities, and also hierarchies of objects. Coverage is particularly weak for these other kinds of knowledge:

1. Time
2. Space
3. Events and actions
4. Cause
5. Collectives
6. Likelihood
7. Obligation
8. Possibility
9. Negation
10. Quantification
11. Continuity and discreteness

6. MODELS OF THE READER

Text prepared without considering the reader is uniformly awful. Programs must have explicit models of the reader that encode (or make available) at least the following four kinds of information:

1. What is obvious--including common factual knowledge and certain "obvious" inferable information. Obviousness does not agree with logical validity.
2. What has already been told, and what is obvious from that.
3. What others believe--including mutual beliefs and beliefs about the writer's belief.
4. What is currently in the reader's attention.

Beyond these, the program should be able to reason about belief and intent.

7. MODELS OF DISCOURSE

This is another linguistic matter, distinct from grammar. Running text has subtle interactions between its parts. When we generate multisentential text, we need a set of principles for organizing it. A few linguists and philosophers are making important contributions, but far more work is needed. Again, to develop effective models of discourse, research projects will need to have linguists on the staff.

An adequate discourse model will include some representation of at least the following:

1. The structures that can be built out of sentences and larger units.
2. The needs of the writer that each discourse structure meets.
3. The principal effects that the use of each structure produces.
4. The effects of various discourse structures on the reader's attention.

8. RELATING THE BASIC COMPONENTS TO TEXT GENERATION

How are these basic components related to the whole task? Why are they necessary, and how does their quality affect what can be done?

8.1. The Text Quality Limitations of Grammars

In order to generate text that is not awkward or misleading, one must be able to control a wide variety of language effects at the sentence level. Effects will be present in the mind of the reader in any case, and so the program must either control them or take serious risks of misunderstanding and error. The effects are produced by the arrangements of words used, and so a theory of the arrangements of words is needed in order to achieve control. Theories of the arrangements of words are (or include) grammars. The ability of a text generation system to express many different ideas well will be ***limited to the different effects controllable through its grammar.***

Use of an ad hoc grammar limits the generator to expressing a narrow range of ideas. It may do well in a short, carefully planned demonstration, but it will be too narrow for many practical purposes.

We can think of the grammar as a bottleneck or filter at the ***output*** of the text generator. Only those expressive techniques that the system can control through its grammar will be used.

8.2. The Knowledge Representation

The knowledge representation frameworks in a program limit the range of things that the program can usefully operate upon. Since a text generator must create text out of ***some*** knowledge representation, it is likewise limited.

Limitations on knowledge representations include two important kinds:

1. Presence of abstractions--are the concepts that must be conveyed in text actually symbolized?
2. Ease of access--is there a fast, uniform method for finding the symbols that represent particular concepts?

We can think of these limits as a bottleneck or filter at the ***input*** of the text generator. Only those concepts that pass through the filter will appear in the text.

8.3. Models of the Reader or System User

Generating acceptable text requires that the generator take into account the knowledge of the reader. If this is not done, text quality is so bad that the results may be useless.[2] To take the reader's knowledge into account requires an explicit model of that knowledge.

Of the four kinds of knowledge previously identified in section 6, the most critical for basic text quality are the knowledge of what is obvious and the knowledge of what has already been conveyed.

8.4. Models of Discourse

We know that single sentences are too limited to express some things. Moving to multisentential text necessarily creates discourse, which involves many kinds of effects that programs cannot yet control. For example, putting one sentence after another can be used to express time sequence, deductive necessity, cause, exemplification or other relationships, ***without any words being used to express the relation.***

Creating these effects when they are desired, and avoiding them when they are not, requires explicit models of discourse phenomena. At a higher level, sequences of sentences and paragraphs of a text must be organized in a principled way. This also requires explicit discourse models. Until such models are developed, texts will be awkward and misunderstanding will be common.

---

[2] With canned text, this problem is usually avoided because the writer knows a great deal about the reader's knowledge.

9. DESIGNING IN 1986 FOR PRACTICAL TEXT GENERATION

What sort of practical application of text generation will be possible in five years? We expect the designer to be in the following situation:

- There will be several examinable systems with developed methods for creating the four basic components. For each kind of component, there will be some attractive precedents for future work. No one system will have a thoroughly elaborated approach to all four.

- System design based on adaptation of these precedents will be possible. The design work will involve creating "handcrafted" systems that embody and reconcile the good techniques. It will require the personal attention of computer scientists, linguists, and programmers who have been involved in the prior research.

- The resulting system can be expected to create acceptable, effective texts, limited by quality considerations to be about one page in length.

For the message-system problem used as a focal problem for the workshop, there were two tasks identified for text generation: a task of reporting system status and a task of reporting how and to whom particular messages would be relevant for an identified collection of people. For both of these tasks it seems feasible for design of a practical text generation module to begin in five years. However, it is questionable whether adequate techniques would be available to ***determine what message relevance to report***.

10. PRESENT RESEARCH STATUS

The most influential research in the next few years will be focused on the four basic components: linguistically justified grammars, knowledge representation formalisms, models of readers, and models of discourse structure and control. Part of the effort will go into developing these components individually, part into learning how to combine them.

Appropriately, most of the current effort is going into either developing single components or combining two of them. Although there are several institutions and individuals working on all four of these components, no one has yet demonstrated a system in which all four approach the scopes of action indicated for them above.

Many important topics are being neglected for lack of research support. (There is no lack of interested people; natural language processing continually generates high interest in the AI community. We are not sure whether there is a shortage of interested ***qualified*** people.)

More information on the state of the art and current activity can be found in the bibliography on text generation, the second part of this report, which includes a section on research in progress.

PART 2. A BIBLIOGRAPHY ON TEXT GENERATION

Contents

1. INTRODUCTION 1

2. BIBLIOGRAPHY 1

3. RESEARCH IN PROGRESS 9

1. INTRODUCTION

This bibliography was prepared in connection with the authors' report on the state of the art in text generation. It includes published works on generation of natural language text by computer programs as well as some prior noncomputational work that has been used as a basis for such computer programs. It is not exhaustive in any sense, and no evaluation is implied by the presence or absence of a citation of any particular publication.

2. BIBLIOGRAPHY

[Allen 78] Allen, J., *Recognizing Intention in Dialogue*. Ph.D. thesis, University of Toronto, 1978.

[Appelt 80a] Appelt, D. E., "A planner for reasoning about knowledge and action," in *Proceedings of the First Annual National Conference on Artificial Intelligence*, Stanford University, August 1980.

[Appelt 80b] Appelt, D. E., "Problem solving applied to language generation," in *Proceedings of the Eighteenth Annual Meeting of the Association for Computational Linguistics*, 1980.

[Appelt 81] Appelt, D., Planning natural language utterances to satisfy multiple goals. Forthcoming Ph.D. thesis, Stanford University.

[Badler 75] Badler, N. I., "The conceptual description of physical activities," in *Proceedings of the Thirteenth Annual Meeting of the Association for Computational Linguistics*, 1975.

[Bates 78] Bates, M., G. Brown, and A. Collins, *Socratic Teaching of Causal Knowledge and Reasoning*, Bolt Beranek and Newman, Inc., Technical Report 3995, December 1978.

[Bates 80] Bates, M., "Language instruction without pre-stored examples," in *Proceedings of the Third Canadian Symposium on Instructional Technology*, February 1980.

[Bates 81a] Bates, M., and R. Ingria, "Controlled transformational sentence generation," in *The Nineteenth Annual Meeting of the Association for Computational Linguistics*, Stanford University, June 1981.

[Bates 81b] Bates, M., et al., "Generative tutorial systems," in *Proceedings of the 1981 ADCIS Conference*, March 1981.

[Bates 81c] Bates, M., et al., *ILIAD Final Report*, Bolt Beranek and Newman, Inc., Technical Report 4771, September 1981.

[Berry 75] Berry, M., *Introduction to Systemic Linguistics: Structures and Systems*, B. T. Batsford, Ltd., London, 1975.

[Berry 77] Berry, M., *Introduction to Systemic Linguistics: Levels and Links*, B. T. Batsford, Ltd., London, 1977.

[Birnbaum 80] Birnbaum, L., M. Flowers, and R. McGuire, "Towards an A.I. model of argumentation," in *Proceedings of the First National Conference on Artificial Intelligence*, Stanford University, August 1980.

[Boguraev 79] Boguraev, B. K., *Automatic Resolution of Linguistic Ambiguities*, Computational Laboratory, Cambridge University, England, Technical Report 11, August 1979.

[Bossie 81] Bossie, S., *A Tactical Component for Text Generation: Sentence Generation Using a Functional Grammar*, University of Pennsylvania. Technical Report MS-CIS-81-5, 1981.

[Brown 74a] Brown, G. P., *Some Problems in German to English Machine Translation*, Massachusetts Institute of Technology, Technical Report 142, December 1974. Project MAC.

[Brown 74b] Brown, R. H., Use of multiple-body interrupts in discourse generation. Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Bachelors Degree thesis.

[Bruce 75] Bruce, B. C., "Generation as a social action," in *Proceedings of Theoretical Issues in Natural Language Processing-I (TINLAP)*, pp. 64-67, Cambridge, Mass., June 1975.

[Bruce 78] Bruce, B. C., A. Collins, A. D. Rubin, and D. Gentner, *A Cognitive Science Approach to Writing*, Bolt Beranek and Newman, Inc., Technical Report 89, June 1978.

[Carbonell 73] Carbonell, J. R., and A. M. Collins, "Natural semantics in artificial intelligence," in *Proceedings of the Third International Joint Conference on Artificial Intelligence*, pp. 344-351, Stanford, Calif., 1973.

[Carr 77] Carr, B., and I. Goldstein, *Overlays: A Theory of Modeling for Computer Aided Instruction*, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, Memo 406, February 1977.

[Chafe 77] Chafe, W. L., "Creativity and verbalization and its implications for the nature of stored knowledge," in R. O. Freedle (ed.), *Discourse Processes: Advances in Research and Theory*. Volume 1: *Discourse Production and Comprehension*, pp. 41-55, Ablex, New Jersey, 1977.

[Chafe 79] Chafe, W. L., "The flow of thought and the flow of language," in T. Givon (ed.), *Syntax and Semantics*. Volume 12: *Discourse and Syntax*, Academic Press, New York, 1979.

[Chester 76] Chester, D., "The translation of formal proofs into English," *Artificial Intelligence* 7, (3), Fall 1976.

[Clancey 78a] Clancey, W. J., *Tutoring Rules for Guiding a Case Method Dialogue*, Stanford University, Department of Computer Science, Heuristic Programming Project, Technical Report HPP-78-25, December 1978. Also appeared in *International Journal of Man-Machine Studies* II, 1979, 25-49.

[Clancey 78b] Clancey, W. J., *An Antibiotic Therapy Selector Which Provides for Explanation*, Stanford University, Technical Report HPP-78-26, December 1978.

[Clancey 79] Clancey, W. J., "Dialogue management for rule-based tutorials," in *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, pp. 155-161, Tokyo, August 1979.

[Clippinger 74] Clippinger, J. H., *A Discourse Speaking Program as a Preliminary Theory of Discourse Behavior and a Limited Theory of Psychoanalytic Discourse*, Ph.D. thesis, University of Pennsylvania, 1974.

[Clippinger 75] Clippinger, J. H., "Speaking with many tongues: Some problems in modelling speakers of actual discourse," in *Proceedings of Theoretical Issues in Natural Language Processing - I (TINLAP)*, pp. 68-73, Cambridge, Mass., June 1975.

[Codd 78] Codd, E. F., et al., *Rendezvous Version 1: An Experimental English-Language Query Formulation System for Casual Users of Relational Databases*, IBM Research Laboratory, San Jose, Calif., Technical Report RJ2144(29407), 1978.

[Cohen 77] Cohen, P. R., and C. R. Perrault, "Overview of 'planning speech acts'," in *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Cambridge, Mass., August 1977.

[Cohen 78] Cohen, P. R., *On Knowing What to Say: Planning Speech Acts*, University of Toronto, Technical Report 118, 1978.

[Cohen 79] Cohen, P. R., and C. R. Perrault, "Elements of a plan-based theory of speech acts," *Cognitive Science* 3, 1979.

[Collins 73] Collins, A. M., J. Passafiume, L. Gould, and J. G. Carbonell, *Improving Interactive Capabilities in Computer-Assisted Instruction*, Bolt Beranek and Newman, Inc., Technical Report 2631, August 1973.

[Cullingford 81] Cullingford, R. E., M. W. Krueger, M. Selfridge, and M. A. Bienkowsky, Automated explanations as a component of a computer-aided design system, 1981. To appear in *IEEE Transactions on Systems, Man, and Cybernetics*.

[Danes 74] Danes, F. (ed.), *Papers on Functional Sentence Perspective*, Academia, Publishing House of the Czechoslovak Academy of Sciences, 1974.

[Davey 79] Davey, A., *Discourse Production*, Edinburgh University Press, Edinburgh, 1979.

[de Beaugrande 80] de Beaugrande, Robert, *Advances in Discourse Processes*. Volume IV: *Text, Discourse, and Process: Toward a Multidisciplinary Science of Texts*, Ablex, Norwood, N. J., 1980.

[de Joia 80] de Joia, A., and A. Stenton, *Terms in Systemic Linguistics*, Batsford Academic and Educational, Ltd., London, 1980.

[Dehn 81a] Dehn, N., "Memory in story invention," in *Proceedings of the Third Annual Conference of the Cognitive Science Society*, University of California, Berkeley, August 1981.

[Dehn 81b] Dehn, N., "Story generation after TALE-SPIN," in *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, University of British Columbia, August 1981.

[Fawcett 80] Fawcett, R. P., *Exeter Linguistic Studies*. Volume 3: *Cognitive Linguistics and Social Interaction*, Julius Groos Verlag Heidelberg and Exeter University, 1980.

[Fillmore 76] Fillmore, C. J., "The case for Case reopened," in Cole, Peter, and Sadock, Jerrold M. (eds.), *Syntax and Semantics*. Volume 8: *Grammatical Relations*, Academic Press, New York, 1976.

[Forbus 81] Forbus, K., and A. Stevens, *Using Qualitative Simulation to Generate Explanations*, Bolt Beranek and Newman, Inc., Technical Report 4490, March 1981. Also appeared in *Cognitive Science* 3, 1981.

[Friedman 69] Friedman, J., "Directed random generation of sentences," *Communications of the ACM* 12, (6), 1969.

[Gabriel 80] Gabriel, R. P., *An Organization for Programs in Fluid Domains*, Ph.D. thesis, Stanford University, 1980.

[Goldman 74] Goldman, N. M., *Computer Generation of Natural Language from a Deep Conceptual Base*, Ph.D. thesis, Stanford University, 1974. Published by Stanford Artificial Intelligence Laboratory as Memo AIM-247.

[Goldman 75a] Goldman, N. M., "The boundaries of language generation," in *Proceedings of Theoretical Issues in Natural Language Processing - I (TINLAP)*, pp. 74-78, Cambridge, Mass., June 1975.

[Goldman 75b] Goldman, N. M., "Conceptual generation," in R. C. Schank (ed.), *Conceptual Information Processing*, North-Holland, Amsterdam, 1975.

[Goldstein 78] Goldstein, I., *Developing a Computational Representation of Problem Solving Skills*, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, Cambridge, Mass., memo 495, October 1978.

[Grimes 75] Grimes, J. E., *The Thread of Discourse*, Mouton, The Hague, 1975.

[Grishman 79] Grishman, R., "Response generation in question-answering systems," in *Proceedings of the Seventeenth Annual Meeting of the Association for Computational Linguistics*, pp. 99-101, August 1979.

[Grosz 79] Grosz, B. J., "Utterance and objective: Issues in natural language communication," in *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, 1979.

[Grosz 80] Grosz, B. J., "Focusing and description in natural language dialogs," in A. Joshi, et al. (eds.), *Elements of Discourse Understanding: Proceedings of a Workshop on Computational Aspects of Linguistic Structure and Discourse Setting*, Cambridge University Press, Cambridge, 1980.

[Habel 77] Habel, C., A. Schmit, and H. Schweppe, *On Automatic Paraphrasing of Natural Language Expressions*, Technische Universiteit, Berlin, Technical Report 3//17, 1977. Semantic Network Project.

[Halliday 61] Halliday, M. A. K., "Categories of the theory of grammar," *Word* 17, 1961.

[Halliday 67a] Halliday, M. A. K., "Notes on transitivity and theme in English," *Journal of Linguistics* 3, (1), 1967, 37-81.

[Halliday 67b] Halliday, M. A. K., "Notes on transitivity and theme in English," *Journal of Linguistics* 3, (2), 1967, 199-244.

[Halliday 68] Halliday, M. A. K., "Notes on transitivity and theme in English," *Journal of Linguistics* 4, (2), 1968, 179-215.

[Halliday 70] Halliday, M. A. K., "Language structure and language function," in J. Lyons (ed.), *New Horizons in Linguistics*, Penguin, 1970.

[Halliday 76a] Halliday, M. A. K., and R. Hasan, *Cohesion in English*, Longman, London, 1976. English Language Series, Title No. 9.

[Halliday 76b] Halliday, M. A. K., *System and Function in Language*, Oxford University Press, London, 1976.

[Halliday 78] Halliday, M. A. K., *Language as Social Semiotic*, University Park Press, Baltimore, Maryland, 1978.

[Heidorn 72] Heidorn, G., *Natural Language Inputs to a Simulation Programming System*, Naval Postgraduate School, Monterey, Calif., Technical Report NPS-55HD72101A, 1972.

[Heidorn 75] Heidorn, G., "Augmented phrase structure grammar," in *Proceedings of Theoretical Issues in Natural Language Processing-I (TINLAP)*, pp. 1-5, Cambridge, Mass., June 1975.

[Herskovits 73] Herskovits, A., *The Generation of French from Semantic Structure*, Stanford Artificial Intelligence Laboratory, Technical Report 212, 1973.

[Hobbs 80] Hobbs, J., and D. Evans, "Conversation as planned behavior," *Cognitive Science* 4, 1980, 349-377.

[Hudson 71] Hudson, R. A., *North-Holland Linguistic Series*. Volume 4: *English Complex Sentences*, North-Holland, London and Amsterdam, 1971.

[Hudson 76] Hudson, R. A., *Arguments for a Non-Transformational Grammar*, University of Chicago Press, Chicago, 1976.

[Hutchins 71] Hutchins, W. J., *The Generation of Syntactic Structures from a Syntactic Base*, North-Holland, Amsterdam, 1971.

[Kay 79] Kay, M., "Functional grammar," in *Proceedings of the Fifth Annual Meeting of the Berkeley Linguistic Society*, 1979.

[Kempen 77] Kempen, G., Building a psychologically plausible sentence generator, 1977. Presented at the Conference on Empirical and Methodological Foundations of Semantic Theories for Natural Language, Nijmegen, The Netherlands.

[Kempen 79] Kempen, G., and E. Hoenkamp, *A Procedural Grammar for Sentence Production*, University of Nijmegen, Department of Psychology, The Netherlands, Technical Report, 1979.

[Klein 65] Klein, S., "Automatic paraphrasing in essay format," *Mechanical Translation* 8, (3), 1965.

[Klein 75] Klein, S., "Meta-compiling text grammars as a model for human behavior," in *Proceedings of Theoretical Issues in Natural Language Processing - I (TINLAP)*, pp. 94-98. Cambridge, Mass., June 1975.

[Knaus 75] Knaus, R., "Incremental sentence processing," *American Journal of Computational Linguistics* Fiche 33, 1975.

[Kripke 77] Kripke, S., "Speaker reference and semantic reference," in P. A. French, et al. (eds.), *Contemporary Perspectives in the Philosophy of Language*, University of Minnesota Press, Minneapolis, 1977.

[Levin 78] Levin, J. A., and N. M. Goldman, *Process Models of Reference in Context*, USC/Information Sciences Institute, RR-78-72, 1978.

[Levy 79a] Levy, D. M., "Communicative goals and strategies: Between discourse and syntax," in T. Givon (ed.), *Syntax and Semantics*, Volume 12: *Discourse and Syntax*, Academic Press, New York, 1979.

[Levy 79b] Levy, D. M., *The Architecture of the Text*, Ph.D. thesis, Stanford University, Department of Computer Science, 1979.

[Linde 75] Linde, C., and W. Labov, "Spatial networks as a site for the study of language and thought," *Language 50* IV, 1975, 924-939.

[Mann 80] Mann, W. C., and J. A. Moore, *Computer as Author--Results and Prospects*, USC/Information Sciences Institute, RR-79-82, 1980.

[Mann 81a] Mann, W. C. and J. A. Moore, "Computer generation of multiparagraph English text," *American Journal of Computational Linguistics* 7, (1), January - March 1981.

[Mann 81b] Mann, W. C., "Two discourse generators," in *The Nineteenth Annual Meeting of the Association for Computational Linguistics*, Sperry Univac, 1981.

[Mann 81c] Mann, W. C., et al., The state of the art of text generation, 1981. To appear in *American Journal of Computational Linguistics*.

[Mann 81d] Mann, W. C., et al., A bibliography on text generation, 1981. To appear in *American Journal of Computational Linguistics*.

[Matthiessen 81] Matthiessen, C. M. I. M., "A grammar and a lexicon for a text-production system," in *The Nineteenth Annual Meeting of the Association for Computational Linguistics*, Sperry Univac, 1981.

[McCoy 81] McCoy, K. F., *Automatic Enhancement of a Database Knowledge Representation for Natural Language Generation*, University of Pennsylvania, Technical Report MS-CIS-81-6, 1981.

[McDonald 75a] McDonald, D. D., "A preliminary report on a program for generating natural language," in *Proceedings of the Third International Joint Conference on Artificial Intelligence*, Tibilisi, USSR, August 1975.

[McDonald 75b] McDonald, D. D., "A framework for writing generation grammars for interactive computer programs," *American Journal of Computational Linguistics* Fiche 33, 1975.

[McDonald 77] McDonald, D. D., "Language generation: The linguistics component (short note)," in *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Cambridge, Mass., August 1977.

[McDonald 78] McDonald, D. D., "Subsequent references: Syntactic and rhetorical constraints," in *Theoretical Issues in Natural Language Processing - 2 (TINLAP)*, ACM, New York, 1978.

[McDonald 80a] McDonald, D. D., *Natural Language Production as a Process of Decision-Making Under Constraints*, Ph.D. thesis, Massachusetts Institute of Technology, Dept. of Electricial Engineering and Computer Science, 1980. To appear as a technical report from the MIT Artificial Intelligence Laboratory.

[McDonald 80b] McDonald, D. D., "The role of discourse structure in language production," in *The Proceedings of the Third Biannual Meeting of the SCSIO/SCEIO*, 1980.

[McDonald 81] McDonald, D. D., "Language production: The source of the dictionary," in *The Nineteenth Annual Meeting of the Association for Computational Linguistics*, Stanford University, June 1981.

[McKeown 79] McKeown, K. R., Paraphrasing using given and new information in a question-answer system, Master's thesis, University of Pennsylvania, Philadelphia, 1979. Number MS-CIS-80-13. Also appeared in *Proceedings of the Seventeenth Annual Meeting of the Association for Computational Linguistics*, August 1979, pp. 67-72.

[McKeown 80a] McKeown, K. R., *Generating Descriptions and Explanations: Applications to Questions about Database Structure*, University of Pennsylvania, Technical Report MS-CIS-80-9, 1980.

[McKeown 80b] McKeown, K. R., "Generating relevant explanations: Natural language responses to questions about database structure," in *Proceedings of The First Annual National Conference on Artificial Intelligence*, pp. 306-309, Stanford, Calif., August 1980.

[McKeown 81] McKeown, K. R., *Generating Natural Language: Deciding What to Say Next*, University of Pennsylvania, Technical Report MS-CIS-81-1, 1981.

[Meehan 75] Meehan, J. R., "Using planning structures to generate stories," *American Journal of Computational Linguistics* Fiche 33, 1975.

[Meehan 77] Meehan, J. R., "TALE-SPIN, an interactive program that writes stories," in *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, August 1977.

[Moore 80] Moore, R., *Reasoning about Knowledge and Action*, SRI International, Artificial Intelligence Center, Technical Note 191, 1980.

[Perrault 78] Perrault, C. R., and P. R. Cohen, *Planning Speech Acts*, University of Toronto, Department of Computer Science, Technical Report, 1978.

[Sacerdoti 77] Sacerdoti, E., *A Structure for Plans and Behavior*, Elsevier North-Holland, Amsterdam, 1977.

[Schank 75] Schank, R., N. Goldman, and C. Reiger, "Inference and paraphrase by computer," *Journal of the ACM* 22, (3), July 1975, 309-328.

[Schlesinger 77] Schlesinger, I. M., *Production and Comprehension of Utterances*. Lawrence Erlbaum Associates, 1977.

[Shapiro 75] Shapiro, S. C., "Generation as parsing from a network into a linear string," *American Journal of Computational Linguistics* Fiche 33, 1975.

[Shapiro 79] Shapiro, S. C., "Generalized augmented transition network grammars for generation from semantic networks," in *Proceedings of the Seventeenth Meeting of the Association for Computational Linguistics*, pp. 25-29, August 1979.

[Simmons 72] Simmons, R., and J. Slocum, "Generating English discourse from semantic networks," *Communications of the ACM* 15, (10), October 1972, 891-905.

[Sleeman 79] Sleeman, D. J., and R. J. Hendley, "ACE: A system which analyses complex explanations," *International Journal of Man-Machine Studies II*, 1979, 125-144.

[Slocum 73] Slocum, J., *Question Answering via Cannonical Verbs and Semantic Models: Generating English from the Model*. University of Texas, Department of Computer Sciences, Austin, Technical Report NL 13, 1973.

[Slocum 75] Slocum, J., "Speech generation from semantic nets," *American Journal of Computational Linguistics* Fiche 33, 1975.

[Stevens 81] Stevens, A., and C. Steinberg, *A Typology of Explanations and its Application to Intelligent Computer Aided Instruction*, Bolt Beranek and Newman, Inc., Technical Report 4626, March 1981.

[Swartout 77] Swartout, W. R., *A Digitalis Therapy Advisor with Explanations*, Massachusetts Institute of Technology, Laboratory for Computer Science, Technical Report, February 1977.

[Swartout 81a] Swartout, W. R., *Producing Explanations and Justifications of Expert Consulting Programs*, Massachusetts Institute of Technology, Technical Report MIT/LCS/TR-251, January 1981.

[Swartout 81b] Swartout, W. R., "Explaining and justifying expert consulting programs," in *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, University of British Columbia, August 1981.

[Thompson 77] Thompson, H. S., "Strategy and tactics: A model for language production," in *Papers from the Thirteenth Regional Meeting*, Chicago Linguistic Society, 1977.

[Thompson 80] Thompson, H. S., *Stress and Salience in English: Theory and Practice*, Xerox Palo Alto Research Center, Technical Report CSL-80-8, May 1980.

[Waltz 78] Waltz, D. L., "An English language question answering system for a large relational database," *Communications of the ACM* 21, (7), July 1978.

[Weiner 80] Weiner, J. L., "BLAH, a system which explains its reasoning," *Artificial Intelligence* 15, November 1980, 19-48.

[Winograd 72] Winograd, T., *Understanding Natural Language*. Academic Press, Edinburgh, 1972.

[Wong 75] Wong, H. K. T., *Generating English Sentences from Semantic Structures*, University of Toronto, Department of Computer Science, Technical Report 84, 1975.

[Yngve 60] Yngve, V. H. A., "A model and a hypothesis for language structure," in *Proceedings of the American Philosophical Society*, pp. 444-466, 1960.

[Yngve 62] Yngve, V. H. A., "Random generation of English sentences," in *The 1961 Conference on Machine Translation of Languages and Applied Language Analysis*, Her Majesty's Stationery Office, London, 1962.

3. RESEARCH IN PROGRESS

This section describes research in text generation either currently in progress or recently completed but not yet described comprehensively in any publication. Like the set of references, it is not exhaustive.

Barbara Grosz and Doug Appelt (SRI International)

Barbara Grosz and Doug Appelt are developing a problem-solving approach to the design of text, extending from prior work by Allen, Cohen, and Perrault. A hierarchical planning system called KAMP (Knowledge and Modalities Planner) is being developed, capable of planning actions that affect another agent's knowledge and wants. It includes critic processes that examine the plan globally for interactions between the effects of actions and propose modifications to the plan that will enable the utterance being planned to realize multiple illocutionary acts. KAMP's knowledge representation is based on Moore's possible worlds semantics approach to reasoning about knowledge and action.

David McDonald (University of Massachusetts, Amherst)

David McDonald is the author of MUMBLE, a system that performs utterance construction, grammatical smoothing, and maintenance of linguistic constraints for natural language generation by expert programs. MUMBLE is available to interested researchers in the common dialect of LISP machine LISP and NIL. The author is currently extending Mumble's grammatical power so that it plans word selection in describing visual scenes and also plans the use of certain connectives such as "but," "also," and "thus."

William Mann and Christian Matthiessen (Information Sciences Institute)

William Mann, Christian Matthiessen, and others are developing the Penman system to explore the problems of creating a portable text generation facility useful in multiple knowledge domains. Penman will seek to deliver knowledge (in English) from inside a system that was not designed to have a text generation component.

The linguistic components of Penman are based on Halliday's Systemic Grammar. A large systemic grammar of English has been implemented and is being fitted with semantic parts.

The knowledge representation, which resembles Brachman's early KL-ONE, is being used for both the subject matter of Penman's generation and the text plans by which Penman generates text. The emphasis of the research is on providing fluent English output from an easily controlled source.

Kathleen McKeown (University of Pennsylvania))

Research is being completed on a text generation system that embodies computational solutions to the questions of what to say next and how to organize it effectively. Two mechanisms are used to handle these problems: 1) rhetorical techniques for communication, encoded as schemas, guide the generation process, and 2) a focusing mechanism helps maintain discourse coherence. Schemas define aspects of discourse structure and are associated with explicit discourse purposes. The focusing mechanism aids in the organization of the message by constraining the selection of what to talk about next to that which ties in with the previous discourse in an appropriate way. This work is being done within the framework of a natural language interface to a database system; the completed system will generate responses to questions about database structure.

Steven Bossie, Kathleen McCoy (University of Pennsylvania))

Two systems are being developed at the University of Pennsylvania in conjunction with McKeown's text generation system. One, developed by Kathleen McCoy, automatically enhances a metalevel description of a database for use by McKeown's text generation system. This system generates subclasses of classes in a given generalization hierarchy. It uses information in the database and a set of axioms to create the subclasses and select salient information describing the subclass divisions.

Steven Bossie is developing a system that will take the ordered message created by McKeown's text generation system and translate it into English. Bossie's system uses a functional grammar, based on a formalism defined by Kay [Kay 79], which will allow for the direct encoding of focus constraints in the grammar. Thus, eventually, the system will use the focusing information provided by McKeown's system to select syntactic constructions.

Rod McGuire (Yale)

Working toward his Ph.D. dissertation, Rod McGuire is developing a model of knowledge representation in human memory to account for observed constraints on the content of oral text. In this model, sentences are generated without building syntactic structures. In multisentential text, coherence arises directly from the form of representation in memory and from memory representation traversal algorithms, using a homogeneous representation to cover syntactic structure, rhetorical structure, and text plans.

Madeline Bates, Robert Ingria, and Kirk Wilson (BBN and Boston University)

The ILIAD system is an intelligent CAI system being developed by Madeline Bates, Robert Ingria, Kirk Wilson[3] and others to give instruction and practice in English. The emphasis is not on teaching grammar, but the system needs to have a deep understanding of the syntactic relationships in the sentences used in examples and exercises. For this reason, the heart of the ILIAD system is a sentence generator that is based on the paradigm of transformational grammar.

ILIAD's grammar blends some aspects of standard transformational theory with the extended standard theory. Rules have been developed to generate not only most of the common English structures but also ungrammatical sentences typical of those produced by people with language-delaying handicaps such as deafness. To control the operation of the generator, several layers of control structures have been developed. Constraints and syntactic specifications allow the user and the system to determine the syntactic form of the sentences at a very high level. Although semantic information is currently used only in lexical insertion, a KL-ONE INTERFACE is being designed.

---

[3] of Learning Tools, Inc., Brookline, Mass.